



UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
10/673,755	09/29/2003	Tushar Deepak Chandra	ARC920030059US1	1473
55508	7590	12/27/2006	EXAMINER	
JOSEPH P. CURTIN, L.L.C. 1469 N.W. MORGAN LANE PORTLAND, OR 97229-5291			VAUTROT, DENNIS L	
			ART UNIT	PAPER NUMBER
			2167	
SHORTENED STATUTORY PERIOD OF RESPONSE		MAIL DATE	DELIVERY MODE	
3 MONTHS		12/27/2006	PAPER	

Please find below and/or attached an Office communication concerning this application or proceeding.

If NO period for reply is specified above, the maximum statutory period will apply and will expire 6 MONTHS from the mailing date of this communication.

Office Action Summary

Application No.

10/673,755

Applicant(s)

CHANDRA ET AL.

Examiner

Dennis L. Vautrot

Art Unit

2167

-- The MAILING DATE of this communication appears on the cover sheet with the correspondence address --

Period for Reply

A SHORTENED STATUTORY PERIOD FOR REPLY IS SET TO EXPIRE 3 MONTH(S) OR THIRTY (30) DAYS, WHICHEVER IS LONGER, FROM THE MAILING DATE OF THIS COMMUNICATION.

- Extensions of time may be available under the provisions of 37 CFR 1.136(a). In no event, however, may a reply be timely filed after SIX (6) MONTHS from the mailing date of this communication.
- If NO period for reply is specified above, the maximum statutory period will apply and will expire SIX (6) MONTHS from the mailing date of this communication.
- Failure to reply within the set or extended period for reply will, by statute, cause the application to become ABANDONED (35 U.S.C. § 133). Any reply received by the Office later than three months after the mailing date of this communication, even if timely filed, may reduce any earned patent term adjustment. See 37 CFR 1.704(b).

Status

- 1) ☒ Responsive to communication(s) filed on 5/25/2006.
- 2a) ☒ This action is **FINAL**. 2b) ☐ This action is non-final.
- 3) ☐ Since this application is in condition for allowance except for formal matters, prosecution as to the merits is closed in accordance with the practice under *Ex parte Quayle*, 1935 C.D. 11, 453 O.G. 213.

Disposition of Claims

- 4) ☒ Claim(s) 1-22 is/are pending in the application.
- 4a) Of the above claim(s) _____ is/are withdrawn from consideration.
- 5) ☐ Claim(s) _____ is/are allowed.
- 6) ☒ Claim(s) 1-22 is/are rejected.
- 7) ☐ Claim(s) _____ is/are objected to.
- 8) ☐ Claim(s) _____ are subject to restriction and/or election requirement.

Application Papers

- 9) ☐ The specification is objected to by the Examiner.
- 10) ☒ The drawing(s) filed on 29 September 2003 is/are: a) ☒ accepted or b) ☐ objected to by the Examiner.
Applicant may not request that any objection to the drawing(s) be held in abeyance. See 37 CFR 1.85(a).
Replacement drawing sheet(s) including the correction is required if the drawing(s) is objected to. See 37 CFR 1.121(d).
- 11) ☐ The oath or declaration is objected to by the Examiner. Note the attached Office Action or form PTO-152.

Priority under 35 U.S.C. § 119

- 12) ☐ Acknowledgment is made of a claim for foreign priority under 35 U.S.C. § 119(a)-(d) or (f).
- a) ☐ All b) ☐ Some * c) ☐ None of:
1. ☐ Certified copies of the priority documents have been received.
 2. ☐ Certified copies of the priority documents have been received in Application No. _____.
 3. ☐ Copies of the certified copies of the priority documents have been received in this National Stage application from the International Bureau (PCT Rule 17.2(a)).

* See the attached detailed Office action for a list of the certified copies not received.

Attachment(s)

- | | |
|--|---|
| 1) <input type="checkbox"/> Notice of References Cited (PTO-892) | 4) <input type="checkbox"/> Interview Summary (PTO-413) |
| 2) <input type="checkbox"/> Notice of Draftsperson's Patent Drawing Review (PTO-948) | Paper No(s)/Mail Date. _____ |
| 3) <input type="checkbox"/> Information Disclosure Statement(s) (PTO/SB/08) | 5) <input type="checkbox"/> Notice of Informal Patent Application |
| Paper No(s)/Mail Date _____ | 6) <input type="checkbox"/> Other: _____ |

DETAILED ACTION

Response to Amendment

1. The applicants' amendment, filed 25 May 2006, has been received, entered into the record and considered.
2. As a result of the amendment, claims 1, 2, 9 – 11, 20 and 21 are amended. Claims 1 – 22 are pending in the application.
3. Applicant's arguments filed regarding claims 1 and 12 have been fully considered but they are not persuasive. On page 10, Applicant argues "Watts does not provide a method of allocating storage to a system user or a storage area network that includes generating a plurality of replicas of each identified master storage image prior to at least one server being allocated to the system user." The cited language from the rejection appears to anticipate this claim however – "Once defined, this server image [the master image] can be rapidly replicated and configured using automated tools to build out images for an entire server pool." This clearly shows multiple replications of the image being generated. Also, "once defined" does not imply that the server has been allocated to the system user already, however it does imply that the image has been defined, which can occur before the server is ever allocated to a system user.

The next argument regarding claim 1 is that the software that is to be executed is installed only once regardless of how many servers will eventually execute the software. The examiner respectfully disagrees with this interpretation. Attention is directed to

Art Unit: 2167

paragraph 60 of **Watt**, where the paragraph notes that there is only one repository regardless of the number and types of storage devices. Examiner does not interpret this to mean that the software is installed only once – rather just one repository or storage device is used. According to the instant application's specification, any number of storage devices could be used, which would include a single storage device, as is the case in this embodiment.

Next, the idea that “each server can execute just one instance at a time and each instance can be mounted by only one server” which is stated in paragraph [0063] still does not appear to teach away from the invention. The next line states that “However, there is no limit to the number of instances that can be created from a server class and made available to a server.” Examiner interprets this to mean that while any given server can only run one system, many replicas [instances] can be made available to the server, which is what appears to be in the gist of the instant application's invention.

The last argument regarding that aspect of claim 1 - “the server images that the Examiner asserts the Watt system allows administrator to rapidly replicate and configure are lists of snapshots, not the claimed replicas of each identified master storage image of claim 1” is also respectfully disagreed with. As mentioned earlier, examiner interprets the “server image” mentioned on page 4, in paragraph [0046], line 6 to be that of the master image from the instant application, because it appears to be referring to the lists of snapshots combined into one master image representing the master storage image. To the Examiner, the master server image is clearly what is being replicated.

The final argument made regarding the last element of the first claim is that "Watts does not provide a method of allocating storage to a system user of a storage area network that includes allocating a selected replica of the claimed plurality of replicas of the master storage image to each server allocated to the system user."

Examiner respectfully disagrees. The particular part of the citation "...build out images for an entire server pool" shows that the images are allocated to each server in the entire server pool. The images were replicated so that there could be one for each of the servers, so the Examiner interprets this to be allocating one for each server.

4. Applicant's arguments filed for claims 2 – 8 and 13 - 19 have been fully considered but they are not persuasive. They are essentially the same as the arguments for 1 and 12, and therefore are disagreed with for the same reasoning.

5. Applicant's arguments filed regarding claims 9 and 20 have been fully considered but they are not persuasive. Applicant argues, "Watt does not generate a read-only copy of the read-only data portion of the claimed selected master storage image." However, examiner interprets paragraph [0097] to refer to a read-only copy of the read-only data portion. In addition to the language cited in the argument which points to the copy "a dependent instance contains copies of just those files necessary to boot the server..." further down in the same paragraph it is noted "Because the dependent instance is mostly shared on a remote, read-only file system..." Therefore, it is interpreted that the copy is read-only.

As for the next argument, **Haun** is not relied upon for that portion of the argument in the office action. However, to address the quoted language in the argument regarding **Haun**, which notes that there is a read-only master copy, rather than a read-only copy of the read-only data portion of the selected master storage image. Examiner believes the read-only master copy would include a copy of the read-only data portion as well, and therefore it would appear to be incorporated by that disclosure.

The final argument is based on the last element of claims 9 and 20 – “allocating a separate writeable data volume of the writeable data portion of the selected master storage image to each server allocated to the system user.” Examiner respectfully disagrees with the Applicant. Specifically, the argument says that **Watt** does not disclose or suggest this limitation. However, Examiner used **Hahn** to reject this limitation. In response to applicant's arguments against the references individually, one cannot show nonobviousness by attacking references individually where the rejections are based on combinations of references. See *In re Keller*, 642 F.2d 413, 208 USPQ 871 (CCPA 1981); *In re Merck & Co.*, 800 F.2d 1091, 231 USPQ 375 (Fed. Cir. 1986).

6. Applicant's arguments filed for claims 10, 11, 21, and 22 have been fully considered but they are not persuasive. They are essentially the same as the arguments for 9 and 20 and therefore are disagreed with for the same reasoning.

Drawings

7. The corrections for the drawings have been received and the objection withdrawn.

Specification

8. The objection to the specification is withdrawn in light of the corrections received.

Claim Rejections - 35 USC § 102

9. The following is a quotation of the appropriate paragraphs of 35 U.S.C. 102 that form the basis for the rejections under this section made in this Office action:

A person shall be entitled to a patent unless –

(e) the invention was described in (1) an application for patent, published under section 122(b), by another filed in the United States before the invention by the applicant for patent or (2) a patent granted on an application for patent by another filed in the United States before the invention by the applicant for patent, except that an international application filed under the treaty defined in section 351(a) shall have the effects for purposes of this subsection of an application filed in the United States only if the international application designated the United States and was published under Article 21(2) of such treaty in the English language.

10. Claims 1 and 12 are rejected under 35 U.S.C. 102(e) as being anticipated by **Watt** (US Patent Application Publication 2003/0126202).
11. Regarding claim 1, **Watt** teaches a method of allocating storage to a system user of a storage area network, the storage area network including storage and a plurality of servers accessing the storage (See page 3-4, paragraph [0046] "Repository manager is responsible for securely and efficiently provisioning and managing server images on storage devices within data center. It allows an administrator to install software one

time such that it can be shared by all servers within a server pool, or even by servers in different pools.” Here, the servers are couple to the repository manager or “storage”), the method comprising:

identifying at least one master storage image that is stored in the storage of the storage area network and that will be associated with a system user when a server is allocated to the system user (See page 4, paragraph [0046] “The administrator can pick and choose from the installed software base to create a master server image....” and see page 4, paragraph [0058] “In an alternate embodiment, repository manager can also install and manage instances of a SAN as well as on a server’s local attached storage.”);

generating a plurality of replicas of each identified master storage image prior to at least one server being allocated to the system user (See page 4, paragraph [0046] “Once defined, this server image can be rapidly replicated and configured using automated tools to build out images for an entire server pool.”);

and allocating a selected replica of the plurality of replicas of the master storage image to each server allocated to the system user (See page 4, paragraph [0046] “Once defined, this server image can be rapidly replicated and configured using automated tools to build out images for an entire server pool.”)

12. Regarding claim 12, **Watt** discloses a storage area network, comprising: a plurality of servers coupled to a storage (See page 3-4, paragraph [0046] “Repository manager is responsible for securely and efficiently provisioning and managing server

Art Unit: 2167

images on storage devices within data center. It allows an administrator to install software one time such that it can be shared by all servers within a server pool, or even by servers in different pools.” Here, the servers are couple to the repository manager or “storage”);

and a storage provisioning device coupled to the servers and allocating at least one server and a portion of the storage to a system user, the storage provisioning device identifying at least one master storage image that is stored in the storage and that will be associated with a system user when a server is allocated to the system user (See page 4, paragraph [0046] “The administrator can pick and choose from the installed software base to create a master server image....” and see page 4, paragraph [0058] “In an alternate embodiment, repository manager can also install and manage instances of a SAN as well as on a server’s local attached storage.”);

the storage provisioning device further generating a plurality of replicas of each identified master storage image prior to at least one server being allocated to the system user (See page 4, paragraph [0046] Once defined, this server image can be rapidly replicated and configured using automated tools to build out images for an entire server pool.”);

and allocating a selected replica of the plurality of replicas of the master storage image to each server allocated to the system user (See page 4, paragraph [0046] Once

Art Unit: 2167

defined, this server image can be rapidly replicated and configured using automated tools to build out images for an entire server pool.”)

Claim Rejections - 35 USC § 103

13. The following is a quotation of 35 U.S.C. 103(a) which forms the basis for all obviousness rejections set forth in this Office action:

(a) A patent may not be obtained though the invention is not identically disclosed or described as set forth in section 102 of this title, if the differences between the subject matter sought to be patented and the prior art are such that the subject matter as a whole would have been obvious at the time the invention was made to a person having ordinary skill in the art to which said subject matter pertains. Patentability shall not be negated by the manner in which the invention was made.

14. Claims 2 - 8 and 13-19 are rejected under 35 U.S.C. 103(a) as being unpatentable over **Watt** (US Patent Application Publication 2003/0126202) as applied to claim 1 above, and further in view of **Sheets et al.** (US 6,819,905).

15. Regarding claim 2, **Watt** teaches a method substantially as claimed. Watt fails to teach de-allocating an allocated replica from the system user each time a server is de-allocated from the system user; and assigning the de-allocated replica to a pool of de-allocated replicas. However, **Sheets et al.** teaches de-allocating an allocated replica from the system user each time a server is de-allocated from the system user; and assigning the de-allocated replica to a pool of de-allocated replicas. (See column 18, lines 32-45 “Another way of looking at how the present invention can dynamically provide hosted service across disparate accounts is to view a portion of the servers as being assigned to a pool of a plurality of virtual servers that may be selectively configured to access software and data for a particular administrative group...it

Art Unit: 2167

automatically allocates one of the servers from the pool of virtual servers to that administrative group. Conversely, if the dispatch module determines that an agency group can relinquish one of its servers, that relinquished server would be added to the pool of virtual servers that are available for re allocation to a different administrative group.") It would have been obvious to one with ordinary skill in the art to combine the method as disclosed in **Watt** with the de-allocating method as disclosed in **Sheets et al.** because the two both disclose methods that operate very similarly, but the de-allocation was simply not addressed in **Watt**, but is a natural extension of **Watt**. By moving the de-allocated replica into the pool, the server is now marked available for future use when it is needed. It is for this reason that one of ordinary skill in the art would have been motivated to include de-allocating an allocated replica from the system user each time a server is de-allocated from the system user; and assigning the de-allocated replica to a pool of de-allocated replicas.

16. Regarding claim 3, **Watt** teaches a method substantially as claimed. **Watt** fails to teach the pool of de-allocated replicas is configured to automatically scrub all replicas in the pool of de-allocated replicas asynchronously from de-allocation the step of de-allocation. However, **Sheets et al.** teaches the pool of de-allocated replicas is configured to automatically scrub all replicas in the pool of de-allocated replicas asynchronously from de-allocation the step of de-allocation. (See column 15, lines 8-14 "One of the significant advantages of the present invention is that the process of reconfiguring servers from one administrative group to a second administrative group

Art Unit: 2167

will wipe clean all of the state associated with a particular customer account for the first administrative group from the reallocated server before that server is brought into service as part of the second administrative group.”) It would have been obvious to one with ordinary skill in the art to combine the method of **Watt** with the automatic scrub function of **Sheets et al.** because of the advantage automatically clearing the unique data provides. This provides a layer of security, allowing different users to use the same equipment dynamically without having to worry about improper sharing of secret data. It is for this reason that one of ordinary skill in the art would have been motivated to have the pool of de-allocated replicas configured to automatically scrub all replicas in the pool of de-allocated replicas asynchronously from de-allocation the step of de-allocation.

17. Regarding claim 4, **Watt** teaches a method substantially as claimed. Watt fails to teach the pool of de-allocated replicas is scrubbed when a number of de-allocated replicas assigned to the pool of de-allocated replicas equals a predetermined number. However, **Sheets et al.** teaches the pool of de-allocated replicas is scrubbed when a number of de-allocated replicas assigned to the pool of de-allocated replicas equals a predetermined number. (See column 15, lines 8-14 “One of the significant advantages of the present invention is that the process of reconfiguring servers from one administrative group to a second administrative group will wipe clean all of the state associated with a particular customer account for the first administrative group from the reallocated server before that server is brought into service as part of the second

Art Unit: 2167

administrative group.”) It would have been obvious to one with ordinary skill in the art to combine the method of **Watt** with the automatic scrub function of **Sheets et al.** because of the advantage automatically clearing the unique data provides. This provides a layer of security, allowing different users to use the same equipment dynamically without having to worry about improper sharing of secret data. While not specifically mentioned here, the predetermined number can be considered to be one and the replica is scrubbed by reconfiguring it for use with another user's data. It is for this reason that one of ordinary skill in the art would have been motivated to have the pool of de-allocated replicas is scrubbed when a number of de-allocated replicas assigned to the pool of de-allocated replicas equals a predetermined number.

18. Regarding claim 5, **Watt** teaches a method substantially as claimed. Watt fails to teach the pool of de-allocated replicas is automatically scrubbed by reformatting. However, **Sheets et al.** teaches the pool of de-allocated replicas is automatically scrubbed by reformatting. (See column 15, lines 8-14 “One of the significant advantages of the present invention is that the process of reconfiguring servers from one administrative group to a second administrative group will wipe clean all of the state associated with a particular customer account for the first administrative group from the reallocated server before that server is brought into service as part of the second administrative group.”) It would have been obvious to one with ordinary skill in the art to combine the method of **Watt** with the automatic scrub by reformatting function of **Sheets et al.** because of the advantage automatically clearing the unique data

Art Unit: 2167

provides. This provides a layer of security, allowing different users to use the same equipment dynamically without having to worry about improper sharing of secret data. It is for this reason that one of ordinary skill in the art would have been motivated to have the pool of de-allocated replicas is automatically scrubbed by reformatting.

19. Regarding claim 6, **Watt** additionally discloses each replica is a logical volume. (See page 5, paragraph [0067] "In another embodiment, the SAN routing and volume assignment can be changed by DSAP system 102 thereby affecting the SAN's mapping of the server's SAN connection to a SAN volume." The way the replica is referred to here makes it clear that the replicas are logical volumes.)

20. Regarding claim 7, **Watt** teaches a method substantially as claimed. Watt fails to teach the system user is one of a customer and an application. However, **Sheets et al.** teaches the system user is one of a customer and an application. (See column 15, lines 8-14 "One of the significant advantages of the present invention is that the process of reconfiguring servers from one administrative group to a second administrative group will wipe clean all of the state associated with a particular customer account for the first administrative group from the reallocated server before that server is brought into service as part of the second administrative group.") It would have been obvious to one with ordinary skill in the art to combine the method of **Watt** with the automatic scrub by reformatting function of **Sheets et al.** because of the advantage automatically clearing the unique data provides. This provides a layer of security, allowing different users to

Art Unit: 2167

use the same equipment dynamically without having to worry about improper sharing of secret data. It is also clear that the user is referring to a customer in **Sheets et al.** It is for this reason that one of ordinary skill in the art would have been motivated to have the system user is one of a customer and an application.

21. Regarding claim 8, **Watt** teaches a method substantially as claimed. Watt fails to teach the master storage image is pre-configured with data and state information that is associated with a system user. However, **Sheets et al.** teaches the master storage image is pre-configured with data and state information that is associated with a system user. (See column 15, lines 8-14 "One of the significant advantages of the present invention is that the process of reconfiguring servers from one administrative group to a second administrative group will wipe clean all of the state associated with a particular customer account for the first administrative group from the reallocated server before that server is brought into service as part of the second administrative group.") It would have been obvious to one with ordinary skill in the art to combine the method of replication as disclosed in **Watt** with the data and state information being associated with the system user in the master storage image as occurs in **Sheets et al.** because of the advantage of being able to keep the data separate for each user. Security is enhanced and the unique data can be kept separate. It is for this reason that one of ordinary skill in the art would have been motivated to have the master storage image pre-configured with data and state information that is associated with a system user.

Art Unit: 2167

22. Regarding claim 13-19, **Watt** in view of **Sheets et al.** disclose the storage area network as cited above for claims 2-8 respectively.

23. Claims 9 and 20 are rejected under 35 U.S.C. 103(a) as being unpatentable over **Watt** (US Patent Application Publication 2003/0126202) in view of **Haun et al.** (US 6,751,058).

24. Regarding claim 9, **Watt** teaches a method of allocating storage between system users of a storage area network, the storage area network including storage and a plurality of servers accessing the storage (See page 4, paragraph [0046] "It allows an administrator to install software one time such at it can be shared by all servers within a server pool, or even by servers in different pools"), the method comprising:

identifying at least one master storage image that is stored in the storage of the storage area network and that will be associated with a system user (See page 4, paragraph [0046] "The administrator and pick and choose from the installed software base to create a master server image...." and see page 4, paragraph [0058] "In an alternate embodiment, repository manager can also install and manage instances of a SAN as well as on a server's local attached storage."), a selected master storage image including both a read-only data portion and a writeable data portion (See page 7, paragraph [0097] "The remainder of the image is shared with other dependent instances by referencing the read-only snapshot containing the original files.");

generating a read-only copy of the read-only data portion of the selected master storage image (See page 7, paragraph [0095] "An independent instance contains an actual physical copy of all files in the master image, with the configuration files updated to provide a unique personally....The independent instance is stored on centralized storage and can be run by any available server.");

sharing the read-only data copy of the read-only data portion of the selected master storage image across the plurality of servers (See page 7, paragraph [0097] "The remainder of the image is shared with other dependent instances by referencing the read-only snapshot containing the original files.");

allocating the read-only copy of the read-only data portion of the selected master storage image to each server allocated to the system user (See page 4, paragraph [0046] Once defined, this server image can be rapidly replicated and configured using automated tools to build out images for an entire server pool.")

Watt fails to teach allocating a separate writable data volume of the writable data portion of the selected master storage image to each server allocated to the system user.

Art Unit: 2167

However, **Haun et al.** teaches allocating a separate writable data volume of the writable data portion of the selected master storage image to each server allocated to the system user. (See column 2, lines 50-55 "According to one embodiment of the present invention, a network computer (NC) system maintains a copy of the operating system that cannot be corrupted by ordinary users of the NC system. Additionally, the NC system may preserve user customizations, such as preferences, by maintaining individual, user, storage areas.")

It would have been obvious to one with ordinary skill in the art to provide a separate writeable data area as disclosed in **Haun et al.** with the allocation method as disclosed in **Watt** because separating the writable area keeps the data from being corrupted as well as providing a layer of security. It is for this reason that one of ordinary skill in the art would have been motivated to allocate a separate writable data volume of the writable data portion of the selected master storage image to each server allocated to the system user.

25. Regarding claim 20, **Watt** in view of **Haun et al.** disclose the storage area network as cited above for claim 9.

26. Claims 10 – 11 and 21 - 22 are rejected under 35 U.S.C. 103(a) as being unpatentable over **Watt** (US Patent Application Publication 2003/0126202) in view of

Haun et al. (US 6,751,058) as applied to claim 9 above, and further in view of **Sheets et al.** (US 6,819,905).

27. Regarding claim 10, **Watt and Haun et al.** teach a method substantially as claimed. **Watt and Haun et al.** fail to teach de-allocating the read-only copy of the read-only data portion of the selected master image from the server to which the read-only copy was allocated when the server is de-allocated from the system user; and de-allocating the writable data volume of the writable data portion of the selected master storage image that was allocated to the de-allocated server. However, **Sheets et al.** teaches de-allocating the read-only copy of the read-only data portion of the selected master image from the server to which the read-only copy was allocated when the server is de-allocated from the system user; and de-allocating the writable data volume of the writable data portion of the selected master storage image that was allocated to the de-allocated server. (See column 18, lines 32-45 "Another way of looking at how the present invention can dynamically provide hosted service across disparate accounts is to view a portion of the servers as being assigned to a pool of a plurality of virtual servers that may be selectively configured to access software and data for a particular administrative group...it automatically allocates one of the servers from the pool of virtual servers to that administrative group. Conversely, if the dispatch module determines that an agency group can relinquish one of its servers, that relinquished server would be added to the pool of virtual servers that are available for re allocation to a different administrative group.") It would have been obvious to one with ordinary skill

Art Unit: 2167

in the art to combine the method as disclosed in **Watt and Haun et al.** with the de-allocating method as disclosed in **Sheets et al.** because they disclose methods that operate very similarly, but the de-allocation was simply not addressed in **Watt and Haun et al.**, but is a natural extension of **Watt and Haun et al.** By moving the de-allocated copy into the pool, the server is now marked available for future use when it is needed. It is for this reason that one of ordinary skill in the art would have been motivated to include de-allocating the read-only copy of the read-only data portion of the selected master image from the server to which the read-only copy was allocated when the server is de-allocated from the system user; and de-allocating the writable data volume of the writable data portion of the selected master storage image that was allocated to the de-allocated server.

28. Regarding claim 11, **Watt and Haun et al.** teach a method substantially as claimed. **Watt and Haun et al.** fail to teach de-allocating the writable data volume includes the steps of: assigning the de-allocated writable data volume to a pool of de-allocated writable data volumes; and scrubbing any writable data volumes assigned to the pool of de-allocated writable data volumes asynchronously from the step of de-allocating the writable data volume. However, **Sheets et al.** teaches de-allocating the writable data volume includes the steps of: assigning the de-allocated writable data volume to a pool of de-allocated writable data volumes; and scrubbing any writable data volumes assigned to the pool of de-allocated writable data volumes asynchronously from the step of de-allocating the writable data volume(See column 18, lines 41-45

Art Unit: 2167

"Conversely, if the dispatch module determines that an agency group can relinquish one of its servers, that relinquished server would be added to the pool of virtual servers that are available for re allocation to a different administrative group." And see column 15, lines 8-14 "One of the significant advantages of the present invention is that the process of reconfiguring servers from one administrative group to a second administrative group will wipe clean all of the state associated with a particular customer account for the first administrative group from the reallocated server before that server is brought into service as part of the second administrative group.") It would have been obvious to one with ordinary skill in the art to combine the teachings of **Watt and Haun et al.** with the teachings in **Sheets et al.** by adding the scrubbing method to enhance the security of sharing the data volumes between different users. It is for this reason that one of ordinary skill in the art would have been motivated to have the step of de-allocating the writable data volume include the steps of: assigning the de-allocated writable data volume to a pool of de-allocated writable data volumes; and scrubbing any writable data volumes assigned to the pool of de-allocated writable data volumes asynchronously from the step of de-allocating the writable data volume.

29. Regarding claim 21 and 22, **Watt and Haun et al.** in view of **Sheets et al.** disclose the storage area network as cited above for claims 10 and 11 respectively.

Conclusion

30. **THIS ACTION IS MADE FINAL.** Applicant is reminded of the extension of time policy as set forth in 37 CFR 1.136(a).


A shortened statutory period for reply to this final action is set to expire THREE MONTHS from the mailing date of this action. In the event a first reply is filed within TWO MONTHS of the mailing date of this final action and the advisory action is not mailed until after the end of the THREE-MONTH shortened statutory period, then the shortened statutory period will expire on the date the advisory action is mailed, and any extension fee pursuant to 37 CFR 1.136(a) will be calculated from the mailing date of the advisory action. In no event, however, will the statutory period for reply expire later than SIX MONTHS from the mailing date of this final action.

Any inquiry concerning this communication or earlier communications from the examiner should be directed to Dennis L. Vautrot whose telephone number is 571-272-2184. The examiner can normally be reached on Monday-Friday 9:00-6:00.

If attempts to reach the examiner by telephone are unsuccessful, the examiner's supervisor, John Cottingham can be reached on 571-272-7079. The fax phone number for the organization where this application or proceeding is assigned is 571-273-8300.

Information regarding the status of an application may be obtained from the Patent Application Information Retrieval (PAIR) system. Status information for published applications may be obtained from either Private PAIR or Public PAIR. Status information for unpublished applications is available through Private PAIR only. For more information about the PAIR system, see <http://pair-direct.uspto.gov>. Should you have questions on access to the Private PAIR system, contact the Electronic Business Center (EBC) at 866-217-9197 (toll-free). If you would like assistance from a USPTO Customer Service Representative or access to the automated information system, call 800-786-9199 (IN USA OR CANADA) or 571-272-1000.

Dv
14 December 2006



JOHN COTTINGHAM
SUPERVISORY PATENT EXAMINER
TECHNOLOGY CENTER 2100

W.L.